Equilibration through local information exchange in networks

K. Y. Michael Wong[1] and David Saad[2]
[1] *Department of Physics, The Hong Kong University of Science and Technology, Clear Water Bay, Hong Kong, China. and*
[2] *The Neural Computing Research Group, Aston University, Birmingham B4 7ET, UK.*
(Dated: November 28, 2005)

We study the equilibrium states of energy functions involving a large set of real variables, defined on the links of sparsely connected networks, and interacting at the network nodes, using the cavity or replica methods. When applied to the representative problem of network resource allocation, an efficient distributed algorithm is devised, with simulations showing full agreement with theory. Scaling properties with the network connectivity and the resource availability are found.

PACS numbers: 02.50.-r, 02.70.-c, 89.20.-a

Many theoretically challenging and practically important problems involve interacting variables connected by network structures [1]. Statistical mechanics of disordered systems makes contributions towards the understanding of such systems at two levels. Macroscopically, it describes the typical behaviour of the networks, using established techniques such as the replica method. Microscopically, it analyses the relation between the variables, using techniques such as the cavity method, that give rise to efficient computational algorithms, such as those reminiscent of the belief propagation algorithm in information processing [2, 3].

Most studies so far have focused on cases of discrete variables. On the other hand, networks of continuous variables were much less explored. There have been recent attempts for message passing of continuous variables *localized* on nodes [4], but many typical problems, such as network resource allocation, involve current variables defined on links *between* nodes.

For optimization on networks, the traditional approach is to adopt global optimization techniques, such as linear or quadartic programming [5]. In contrast, message-passing approaches have the potential to solve global optimization problems via local updates, thereby reducing the computational complexity. An even more important advantge, relevant to practical implementation, is its distributive nature. Since it does not require a global optimizer, it is particularly suitable for distributive control in large or evolving networks.

In this paper we study a system with real variables that can be mapped onto a sparse graph and suggest an efficient message passing approximation method. After formulating the problem at a general temperature, we focus on a prototype for optimization, termed resource allocation and well known in the areas of computer science and operations management [6, 7].

We consider a sparse network with N nodes, labelled $i = 1, \ldots, N$. Each node i is randomly connected to c other nodes. The connectivity matrix is given by $\mathcal{A}_{ij} = 1, 0$ for connected and unconnected node pairs respectively. A link variable y_{ij} is defined on each connected link from j to i. We consider an energy function (cost) $E = \sum_{(ij)} \mathcal{A}_{ij}\phi(y_{ij}) + \sum_i \psi(x_i, \{y_{ij}|\mathcal{A}_{ij} = 1\})$, where x_i is a quenched variable defined on node i. In the context of probabilistic inference, y_{ij} may represent the coupling between observables in nodes j and i, $\phi(y_{ij})$ may correspond to the logarithm of the prior distribution of y_{ij}, and $\psi(x_i, \{y_{ij}|\mathcal{A}_{ij} = 1\})$ the logarithm of the likelihood of the observables x_i. In the context of resource allocation, $y_{ij} \equiv -y_{ji}$ may represent the current from node j to i, $\phi(y_{ij})$ may correspond to the transportation cost, and $\psi(x_i, \{y_{ij}|\mathcal{A}_{ij} = 1\})$ the performance cost of the allocation task on node i, dependent on the node capacity x_i.

We are interested in the case of intensive connectivity $c \sim O(1) \ll N$. Since the probability of finding a loop of finite length on the nework is low, the cavity method well describes the local environment of a node. A node is connected to c branches in a tree structure, and the correlations among the branches of the tree are neglected. In each branch, nodes are arranged in generations. A node is connected to an ancestor node of the previous generation, and another $c-1$ descendent nodes of the next generation. Considering node i as the ancestor of node j, the descendents of node j form a tree structure $\mathbf{T}$ with $c-1$ branches, labelled by $k \neq i$ for $\mathcal{A}_{jk} = 1$. At a temperature $T \equiv \beta^{-1}$, the free energy $F(y_{ij}|\mathbf{T})$ can be expressed in terms of the free energies $F(y_{jk}|\mathbf{T}_k)$ of its descendents. The free energy can be considered as the sum of two parts, $F(y|\mathbf{T}) = N_{\mathbf{T}}F_{\rm av} + F_V(y|\mathbf{T})$, where $N_{\mathbf{T}}$ is the number of nodes in the tree $\mathbf{T}$, $F_{\rm av}$ is the average free energy per node, and $F_V(y|\mathbf{T})$ is referred to as the *vertex free energy*. This leads to the recursion relation

$$F_V(y_{ij}|\mathbf{T}) = -T\ln\left\{\prod_{k\neq i}\left(\int dy_{jk}\right)\exp\left[-\beta\psi(x_j,\{y_{jk}\}) - \beta\sum_{k\neq i}\left(F_V(y_{jk}|\mathbf{T}_k)+\phi(y_{jk})\right)\right]\right\}\Bigg|_{\mathcal{A}_{jk}=1} - F_{\rm av}, \quad (1)$$

$$F_{\rm av} = -T\left\langle \ln\left\{\prod_{k}\left(\int dy_{jk}\right)\exp\left[-\beta\psi(x_j,\{y_{jk}\}) - \beta\sum_{k}\left(F_V(y_{jk}|\mathbf{T}_k)+\phi(y_{jk})\right)\right]\right\}\Bigg|_{\mathcal{A}_{jk}=1}\right\rangle_x, \quad (2)$$

where $\mathbf{T}_k$ is the tree terminated at node k, and $\langle \ldots \rangle_x$

represents the average over the distribution of x.

For more concrete discussions, we focus on the resource allocation problem, which is applicable to typical situations where a large number of nodes are required to balance loads/resources, such as reducing internet traffic congestion and streamlining network flows of commodities [8]. In computer science, many practical solutions are usually heuristic and focus on practical aspects (e.g., communication protocols). Here we study a more generic version of the problem. In the context of computer networks, it is represented by nodes of some computational power that should carry out tasks. Both computational powers and tasks will be chosen at random from some arbitrary distribution. The nodes are located on a randomly chosen sparse network of some connectivity. The goal is to allocate tasks on the network such that demands will be satisfied while the migration of (sub-)tasks is minimised.

We focus here on the satisfiable case where the total computing power is greater than the demand, and where the number of nodes involved is very large. This is of interest to physicists due to the applicability of the techniques we introduce to the analysis of sparsely connected systems with real variables. Each node on the network has a capacity (computational capability minus allocated tasks) x_i randomly drawn from a distribution $\rho(x_i)$. (The algorithms presented later can easily accommodate any connectivity profile within the same framework.) With the aim to satisfy the capacity constraints, we have $\psi(x_i, \{y_{ij}|\mathcal{A}_{ij}=1\}) = \ln[\Theta(\sum_j \mathcal{A}_{ij}y_{ij} + x_i) + \epsilon]$, where $\epsilon \to 0$. The problem reduces to the load balancing task of minimizing the energy function (cost) $E = \sum_{(ij)} \mathcal{A}_{ij}\phi(y_{ij})$, subject to the capacity constraints.

When $\phi(y)$ is a general even function of the current y, we may also derive Eq. (1) using the replica method. We first introduce the *chemical potentials* μ_i of nodes i, and approximate the current y_{ij} as driven by the potential differences between nodes $y_{ij} = \mu_j - \mu_i$. Since sparse networks are locally tree-like, the probability of finding short loops is vanishing in large networks, and the approximation works well.

Considering the optimization problem in the space of chemical potentials, we calculate the replicated partition function $\langle Z^n \rangle_{\mathcal{A},x}$ averaged over the connectivity matrix and the capacity distribution, and take the limit $n \to 0$. Assuming replica symmetry, the saddle point equations of the replica method yields a recursion relation for a two-component function R dependent on the tree structure $\mathbf{T}$, given by

$$R(z,\mu|\mathbf{T}) = \frac{1}{\mathcal{D}} \prod_{k=1}^{c-1} \left(\int d\mu_k R(\mu,\mu_k|\mathbf{T}_k) \right) \Theta\left(\sum_{k=1}^{c-1} \mu_k - c\mu + z + x_{V(\mathbf{T})} \right) \exp\left(-\frac{\beta\epsilon}{2}\mu^2 - \beta \sum_{k=1}^{c-1} \phi(\mu-\mu_k) \right), \quad (3)$$

where $\mathcal{D}$ is a constant, $\mathbf{T}_k$ represents the tree terminated at the k^{th} descendent, and $x_{V(\mathbf{T})}$ is the capacity of the vertex of the tree $\mathbf{T}$. The term $\beta\epsilon\mu^2/2$, with $\epsilon \to 0$, is introduced to break the translational symmetry of the chemical potentials, since the energy function is invariant under the addition of an arbitrary global constant to all chemical potentials.

Eq. (3) expresses $R(z,\mu|\mathbf{T})$ in terms of $c-1$ functions $R(\mu,\mu_k|\mathbf{T}_k)$ $(k = 1,..,c-1)$, a characteristic of the tree structure. Furthermore, except for the factor $\exp(-\beta\epsilon\mu^2/2)$, R is a function of $y \equiv \mu - z$, which is interpreted as the current drawn from a node with chemical potential μ by its ancestor with chemical potential z. One can then express the function R as the product of a *vertex partition function* Z_V and a normalization factor W, that is, $R(z,\mu|\mathbf{T}) = W(\mu)Z_V(y|\mathbf{T})$. In the limit $n \to 0$, the dependence on μ and y decouples, enabling one to derive a recursion relation for the *vertex free energy* $F_V(y|\mathbf{T}) \equiv -T \ln Z_V(y|\mathbf{T})$ and arrive at Eq. (3).

The current distribution and the average free energy per link can be derived by integrating the current y' in a link from one vertex to another, fed by the trees $\mathbf{T}_1$ and $\mathbf{T}_2$, respectively; the obtained expressions are $P(y) = \langle \delta(y - y') \rangle_\star$ and $\langle E \rangle = \langle \phi(y') \rangle_\star$ where $\langle \bullet \rangle_\star = \left\langle \int dy' \exp\left[-\beta E(y')\right] (\bullet) / \int dy' \exp\left[-\beta E(y')\right] \right\rangle_x$, and $E(y') = F_V(y'|\mathbf{T}_1) + F_V(-y'|\mathbf{T}_2) + \phi(y')$.

The solution of Eq. (1) can be obtained numerically for optimization ($T = 0$). Since the vertex free energy of a node depends on its own capacity and the disordered configuration of its descendents, we generate 1000 nodes at each iteration of Eq. (1), with capacities randomly drawn from the distribution $\rho(x)$, and each being fed by $c-1$ nodes randomly drawn from the previous iteration.

We have discretized the vertex free energy function into a vector, whose i^{th} component is the value $F_V(y_i|\mathbf{T})$. To speed up the optimization search at each node, we first find the *vertex saturation current* drawn from a node such that: (a) the capacity of the node is just used up; (b) the current drawn by each of its descendant nodes is just enough to saturate its own capacity constraint. At this saturation point, we can separately optimize the current drawn by each descendant node, providing a convenient starting point for searching the optimal solutions.

Figure 1(a) shows the results of iteration for a Gaussian capacity distribution $\rho(x)$ with variance 1 and average $\langle x \rangle$. Each iteration corresponds to adding one extra generation to the tree structure, such that the iterative process corresponds to approximating the network by an increasingly extensive tree. We observe that after an initial rise with iteration steps, the average energies converge to steady-state values, at a rate which increases with the average capacity.

To study the convergence rate of the iterations, we fit the average energy at iteration step t using $\langle E(t) - E(\infty) \rangle \sim \exp(-\gamma t)$ in the asymptotic regime. As shown in the inset of Fig. 1(a), the relaxation rate γ increases with the average capacity. A cusp appears at the average capacity of about 0.45, below which convergence is slow due to a plateau that develops in the average energy curve before the final stage. The slowdown is prob-

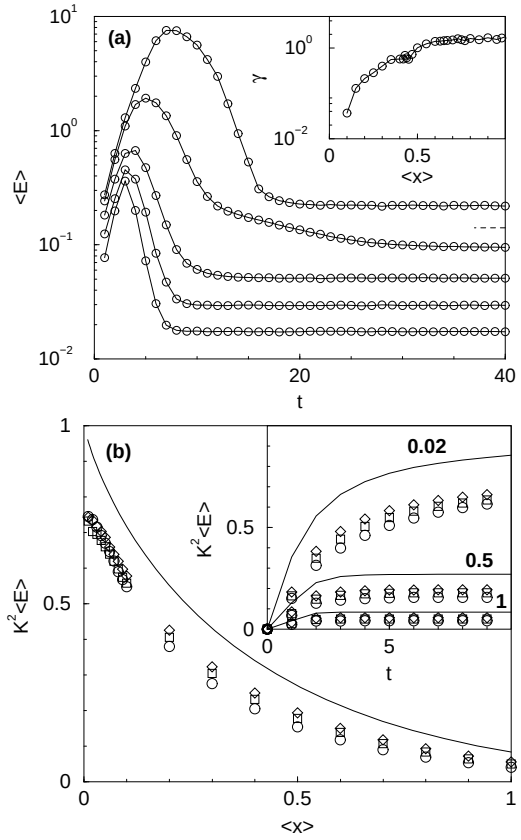

FIG. 1: Results for $N = 1000$ and $\phi(y) = y^2/2$. (a) $\langle E\rangle$ obtained by iterating Eq. (1) as a function of t for $\langle x\rangle = 0.1$, 0.2, 0.4, 0.6, 0.8 (top to bottom) and $c=3$. Dashed line: The asymptotic $\langle E\rangle$ for $\langle x\rangle = 0.1$. Inset: γ as a function of $\langle x\rangle$. (b) $K^2\langle E\rangle$ as a function of $\langle x\rangle$ for $c = 3$ (◯), 4 (□), 5 (◇), large K (line). Inset: $K^2\langle E\rangle$ as a function of time for random sequential update of Eqs. (4-6). Symbols: same as (a).

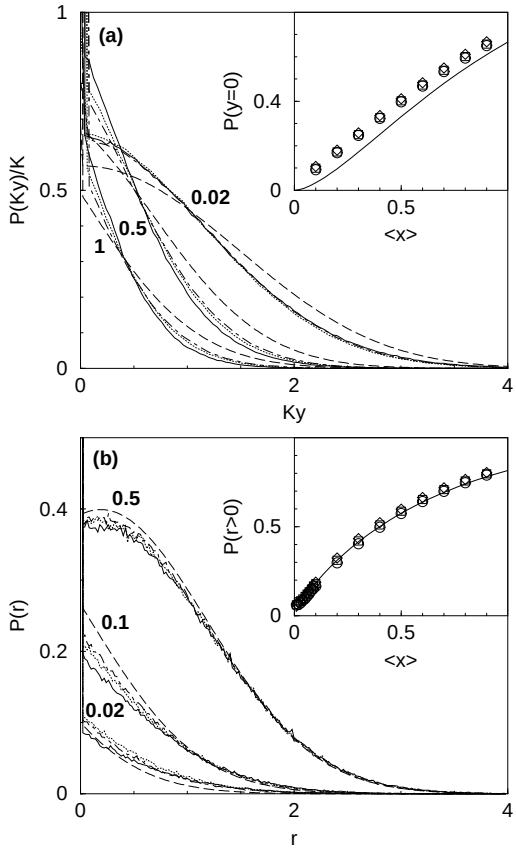

FIG. 2: Results for $N = 1000$ and $\phi(y) = y^2/2$. (a) The current distribution $P(Ky)/K$ for $\langle x\rangle = 0.02, 0.5, 1$, and $c = 3$ (solid lines), 4 (dotted lines), 5 (dot-dashed lines), large K (long dashed lines). Inset: $P(y=0)$ as a function of $\langle x\rangle$ for $c = 3$ (◯), 4 (□), 5 (◇), large K (line). (b) The resource distribution $P(r)$ for $\langle x\rangle = 0.02, 0.1, 0.5$, large K. Symbols: same as (a). Inset: $P(r>0)$ as a function of $\langle x\rangle$. Symbols: same as (a) inset.

ably due to the appearance of increasingly large clusters of nodes with negative capacities, which draw currents from increasingly extensive regions of nodes with excess capacities to satisfy the demand.

The local nature of the recursion relation (1) points to the possibility that the network optimization can be solved by message passing approaches, However, in contrast to other message passing algorithms which pass conditional probability estimates of *discrete values* to the neighboring nodes, the messages in the present context are more complex, since they are *functions* $F_V(y|\mathbf{T})$ of the current y. We simplify the message to 2 parameters, namely, the first and second derivatives of the vertex free energies. For the quadratic load balancing task, it can be shown that a self-consistent solution of the recursion relation Eq. (1) consists of vertex free energies which are piecewise quadratic with continuous slopes. This makes the 2-parameter message a very precise approximation.

Let $(A_{ij}, B_{ij}) \equiv (\partial F_V(y_{ij}|\mathbf{T}_j)/\partial y_{ij}, \partial^2 F_V(y_{ij}|\mathbf{T}_j)/\partial y_{ij}^2)$ be the message passed from node j to i. Using Eq. (1), the recursion relations lead to the forward message (A_{ij}, B_{ij}) from node j to i, followed by the backward message y_{jk} from node j to k,

$$A_{ij} \leftarrow -\mu_{ij}, \quad B_{ij} \leftarrow \frac{\Theta(-\mu_{ij}+\epsilon)}{\sum_{k\neq i}\mathcal{A}_{jk}(\phi''_{jk}+B_{jk})^{-1}}, \tag{4}$$

$$y_{jk} \leftarrow y_{jk} - \frac{\phi'_{jk}+A_{jk}+\mu_{ij}}{\phi''_{jk}+B_{jk}}, \tag{5}$$

$$\mu_{ij} = \min\Bigg\{\Bigg[\sum_{k\neq i}\mathcal{A}_{jk}[y_{jk}-(\phi'_{jk}+A_{jk})(\phi''_{jk}+B_{jk})^{-1}] + \Lambda_j - y_{ij}\Bigg]\Bigg[\sum_{k\neq i}\mathcal{A}_{jk}(\phi''_{jk}+B_{jk})^{-1}\Bigg]^{-1}, 0\Bigg\}, \tag{6}$$

with ϕ'_{jk} and ϕ''_{jk} representing the first and second derivatives of $\phi(y)$ at $y = y_{jk}$ respectively.

We note that Eqs. (4-6) differ from conventional message-passing algorithms in that backward messages of the currents are present. As a consequence of representing the messages by the first and second derivatives,

the backward messages serve to inform the descendent nodes the particular arguments they should use in calculating the derivatives for sending the next messages. Furthermore, the criterion that $y_{ij} = -y_{ji}$ provides a check for the convergence of the algorithm.

Another usage of the backward messages is in monitoring the optimal cost function during simulations. This saves the extra step of calculating the energy associated with a link in the conventional Bethe approach.

For the quadratic load balancing task considered here, an independent exact optimization is available for comparison. The chemical potentials turn out to be the Lagrange multipliers of the capacity constraints, and the relation between the currents and the chemical potentials turn out to be exact. The Kühn-Tucker conditions for the optimal solution yield

$$\mu_i = \min\left[\frac{1}{c}\left(\sum_j \mathcal{A}_{ij}\mu_j + \Lambda_i\right), 0\right]. \qquad (7)$$

Like in the message-passing algorithm, this condition also provides a local iterative solution to the optimization problem. Simulations show that it yields excellent agreement with Eqs. (1) and (4-6).

To study the dependence on the connectivity, we first consider the limit of large $K \equiv c - 1$. In this limit, Eq. (4) converges to the steady-state results of $A_{ij} = \max[K^{-1}\sum_{k\neq i}\mathcal{A}_{jk}A_{jk} - x_j), 0]$ and $B_{ij} \sim K^{-1}$. Then $\sum_{k\neq i}\mathcal{A}_{jk}A_{jk}$ becomes self-averaging and equal to Km_A, where $m_A \sim K^{-1}$ is the mean of the messages A_{ij}. Thus, $y_{ij} \sim K^{-1}$, $\mu_i \sim K^{-1}$, and $\langle E\rangle \sim K^{-2}$. The physical picture of this scaling behavior is that the total current required by a node to satisfy its capacity constraint is shared by the links. After rescaling, the physical quantities such as $K^2\langle E\rangle$, $P(Ky)/K$ and $P(K\mu)/K$ become purely dependent on the average capacity $\langle x\rangle$.

For increasing finite values of K, Fig. 1(b) shows the common trend of $K^2\langle E\rangle$ decreasing with $\langle x\rangle$ exponentially, and gradually approaching the large K limit. The scaling property extends to the dynamics of optimization (Fig. 1(b) inset). As shown in Fig. 2(a), the current distribution $P(Ky)/K$ consists of a delta function component at $y = 0$ and a continuous component, whose breadth decreases with $\langle x\rangle$. Remarkably, the distributions for different connectivities collapse almost perfectly after the currents are rescaled by K^{-1}, with a very mild dependence on K and gradually approaching the large K limit. As shown in the inset of Fig. 2(a), the fraction of idle links increases with $\langle x\rangle$. Hence the current-carrying links form a percolating cluster at a low $\langle x\rangle$, and breaks into isolated clusters at a high $\langle x\rangle$. The fraction has a weak dependence on the connectivity, confirming the almost universal distributions rescaled for different K.

Since the current on a link scales as K^{-1}, the allocated resource of a node should have a weak dependence on the connectivity. Defining the resource at node i by $r_i \equiv x_i + \sum_j \mathcal{A}_{ij}y_{ij}$, the resource distribution $P(r)$ shown in Fig. 2(b) confirms this behavior even at low connectivities. The fraction of nodes with unsaturated capacity constraints increases with the average capacity, and is weakly dependent on the connectivity (Fig. 2(b) inset). Hence the saturated nodes form a percolating cluster at a low average capacity, and breaks into isolated clusters at a high average capacity, It is interesting to note that at the average capacity of 0.45, below which a plateau starts to develop in the relaxation rate of the recursion relation, Eq. (1), the fraction of unsaturated nodes is about 0.53, close to the percolation threshold of 0.5 for $c = 3$.

In summary, using the example of the resource allocation problem on sparsely connected networks, we have demonstrated the use of message-passing methods for equilibration. This extends the conventional usage of Bayesian message passing for inference in problems with discrete variables to problems with continuous variables (such as in optimization), opening up a rich area for further investigations with many potential applications. The study also reveals the scaling properties of the resource allocation model, showing that the resource distribution on the nodes depends principally on the networkwide availability of resources, and is only weakly dependent on the connectivity. The links share the task of resource provision among themselves, leading to current distributions that are almost universally dependent on the resource availability after rescaling.

Acknowledgments

This work is partially supported by the Research Grant Council of Hong Kong (grants HKUST6062/02P, DAG04/05.SC25 and DAG05/06.SC36) and EVERGROW, IP No. 1935 in FP6 of the EU.

[1] H. Nishimori, Statistical Physics of Spin Glasses and Information Processing, (OUP, Oxford UK, 2001).
[2] M. Opper and D. Saad, Advanced Mean Field Methods (MIT press, Cambridge, MA, 2001).
[3] D. J. C. Mackay, Information Theory, Inference and Learning Algorithms, (CUP, Cambridge UK, 2003).
[4] N. Skantzos, I. P. Castillo, and J. P. L. Hatchett, arXiv:cond-mat/0508609 (2005).
[5] D. Bertsekas, Linear Network Optimzation (MIT Press, Cambridge, MA, 1991).
[6] L. Peterson and B. S. Davie, Computer Networks: A Systems Approach, (Academic Press, San Diego CA, 2000).
[7] Y. C. Ho, L. Servi, and R. Suri, Large Scale Systems **1**, 51 (1980).
[8] S. Shenker, D. Clark, D. Estrin S. Herzog, ACM Comp. Comm. Rev. **26**. 19 (1996).